Harris, R. D. F., & Nguyen, A. T. H. (2017). Dynamic factor long memory volatility. *Quantitative Finance*, 17(8), 1205-1221.
<https://doi.org/10.1080/14697688.2016.1260757>

Peer reviewed version

Link to published version (if available):
[10.1080/14697688.2016.1260757](https://doi.org/10.1080/14697688.2016.1260757)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Taylor & Francis at <https://www.tandfonline.com/doi/full/10.1080/14697688.2016.1260757> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Dynamic Factor Long Memory Volatility

Richard D.F. Harris^a

University of Exeter

Anh Nguyen^b

Foreign Trade University, Vietnam

November 2016

Abstract

In this paper, we develop a long memory orthogonal factor (LMOF) multivariate volatility model for forecasting the covariance matrix of financial asset returns. We evaluate the LMOF model using the volatility timing framework of Fleming et al. (2001) and compare its performance with that of both a static investment strategy based on the unconditional covariance matrix and a range of dynamic investment strategies based on existing short memory and long memory multivariate conditional volatility models. We show that investors should be willing to pay to switch from the static strategy to a dynamic volatility timing strategy and that, among the dynamic strategies, the LMOF model consistently produces forecasts of the covariance matrix that are economically more useful than those produced by the other multivariate conditional volatility models, both short memory and long memory. Moreover, we show that combining long memory volatility with the factor structure yields better results than employing either long memory volatility or the factor structure alone. The factor structure also significantly reduces transaction costs, thus increasing the feasibility of dynamic volatility timing strategies in practice. Our results are robust to estimation error in expected returns, the choice of risk aversion coefficient, the estimation window length and sub-period analysis.

Keywords: Finance; Conditional variance-covariance matrix; Long memory; Factor models; Volatility timing.

^aAddress: Xfi Centre for Finance and Investment, University of Exeter, Exeter, EX4 4PU, UK. Email: R.D.F.Harris@exeter.ac.uk (corresponding author). ^bAddress: Foreign Trade University, Hochiminh City Campus, 15 D5, Binh Thanh District, Hochiminh City, Vietnam. Email: anhth.nguyen@ftu.edu.vn. We are grateful to the editor and two referees for their comments that have helped to improve the paper.

1 Introduction

Factor models are widely used in asset pricing, asset allocation and risk management to forecast the covariance matrix of returns. The advantages of factor models have been well documented in the literature and empirically confirmed in practice. For example, Chan et al. (1999) study the performance of various factor models in portfolio optimization and show that fundamental factor models clearly improve the forecasts of the covariance matrix. In a similar study, Amenc and Martellini (2002) use an implicit factor model to estimate the covariance matrix for hedge fund index returns and find that portfolios constructed using the factor-based estimator have lower ex post volatility than both value-weighted and equally-weighted portfolios. Recent studies incorporate time-varying conditional volatility in the factor structure and suggest that this can lead to significant economic benefits. For example, Briner and Connor (2008) embed an exponential weighting in the covariance matrix of the factors and show that the conditional factor EWMA model outperforms the fully estimated EWMA model in terms of forecast performance. Han (2006) develops a dynamic factor multivariate stochastic volatility model, which utilises a set of latent factors to capture the dynamics of the first two moments of returns, and shows that it provides a significant performance improvement relative to using the unconditional covariance matrix estimator. Alessi et al. (2009) build a Dynamic Factor GARCH model by combining generalised dynamic factor models with a multivariate GARCH structure and show that their model performs better than a static factor model and a univariate GARCH model in forecasting the conditional variances and covariances of multivariate inflation series and financial asset returns. Using a similar approach, Barigozzi and Hallin (2016) propose a two-step generalised dynamic factor model that accounts for a factor structure in both returns and volatilities. Incorporating GARCH processes in modelling volatilities, they show that their model produces better forecasts of the covariance matrix than those generated from multivariate GARCH and static factor GARCH models.

In dynamic factor volatility models, the volatilities of the factors are typically estimated using established conditional volatility models such as EWMA, GARCH or Stochastic Volatility, in which shocks to volatility dissipate rapidly due to their exponential weighting. However, a growing body of empirical evidence suggests that volatility exhibits longer memory than these models imply (see, for example, Taylor, 1986, Ding et al., 1993, Andersen et al., 2001). This has prompted the development of a number of long memory volatility models, which have been shown to produce superior forecast performance relative to their short memory counterparts (Baillie et al., 1996, Engle and Lee, 1999, Vilasuso, 2002, Andersen et al., 2003, Davidson,

2004, Pong et al., 2008, Corsi, 2009). Though appealing in the univariate setting, multivariate long memory volatility models suffer from a curse-of-dimensionality problem, and this severely constrains their practical application in asset allocation and risk management. Financial practitioners have thus largely eschewed the well-specified long memory volatility models, relying instead on the simpler historical volatility or short memory EWMA models.

In this paper, we propose a dynamic factor long memory conditional volatility model that can be implemented in the context of the high dimensional covariance matrices that are typically encountered in financial applications. The model exploits the simple structure of factor models while capturing the long memory feature of volatility. Its simplicity circumvents the computational burden of multivariate long memory models, making it feasible for high dimensional volatility modelling. The Long Memory Orthogonal Factor (LMOF) volatility model that we propose is achieved by embedding the long memory LM-EWMA model of Zumbach (2006) into an orthogonal factor structure. The LM-EWMA model is a highly parsimonious model that captures the long memory of a process by combining short memory EWMA processes at different time horizons.¹ The simplicity of the LM-EWMA process allows us to adopt a richer specification than is normally assumed in the volatility factor structure. Specifically, we assume that the volatilities of both the factors and the idiosyncratic shocks are time-varying and characterised by long memory.

We quantify the economic value of the LMOF model using the volatility timing framework of Fleming et al. (2001), in which expected returns are assumed to be constant and investors periodically update their portfolios based on forecasts of the conditional covariance matrix. If the covariance matrix is time-invariant, the optimal weights will be constant over time and so an investor would follow a static strategy. However, if the investor believes that the covariance matrix is time-varying, he will follow a dynamic strategy in which the optimal weights are adjusted on the basis of his forecast of the conditional covariance matrix. We compare the economic value of the LMOF model against that of a wide range of multivariate EWMA and GARCH conditional volatility models.² In particular, the benchmark models include two

¹ Harris and Nguyen (2013) provide evidence on the superior forecast performance of two long memory volatility models based on the LM-EWMA model of Zumbach (the multivariate long memory LM-EWMA model of Zumbach (2011) and the univariate LM-EWMA combined with the Dynamic Conditional Correlation (DCC) framework of Engle (2002)), against a wide range of long memory and short memory multivariate volatility models.

² As is common in the literature, we restrict our attention to the class of EWMA and GARCH models. We do not consider multivariate realized volatility models owing to the nature of our data. Nor do we

multivariate long memory models – the multivariate LM-EWMA of Zumbach (2011) and the component CGARCH model of Engle and Lee (1999) implemented using the DCC framework – and two widely used short memory volatility models – the Riskmetrics’ EWMA model of JPMorgan (1994) and the GARCH-DCC model. We apply the models to two datasets: the first is a portfolio of 21 international stock indices and 13 international bond indices over the period 1 January 1988 to 31 December 2013, while the second is a portfolio of individual stocks drawn from the Dow Jones Industrial Average index over the period 1 June 1999 to 31 December 2013. We measure portfolio performance using the out-of-sample Sharpe ratio, the abnormal return and the performance fees that investors would be willing to pay to switch from the static to the dynamic strategies. We also calculate the breakeven transaction costs that make investors indifferent between the static and the dynamic strategies in terms of utility. We report two main findings. First, consistent with the literature, investors should be willing to pay to switch from the static strategy to a dynamic volatility timing strategy at all rebalancing frequencies. Sub-period analysis shows that dynamic strategies outperform the static portfolios more often in market downturns than in normal market conditions. Second, among the dynamic strategies, the factor-based LMOF model generally produces forecasts of the covariance matrix that are economically more useful than those produced by other multivariate conditional volatility models, both short memory and long memory. The factor structure also significantly reduces transaction costs, thus increasing the feasibility of dynamic volatility timing strategies in practice. These results apply to both datasets and are robust to estimation error in expected returns, the choice of risk aversion coefficient and the estimation window length. The results also suggest that combining long memory volatility with the factor structure yields better results than employing either long memory volatility or the factor structure alone.

The remainder of the paper is organised as follows. In Section 2, we describe the Long Memory Orthogonal Factor conditional volatility model. Section 3 presents the volatility timing framework that is used to study the economic benefits of the dynamic strategies. The empirical methodology is described in Section 4. Section 5 reports the empirical results, while Section 6 offers some concluding comments and suggestions for future research.

consider multivariate stochastic volatility models on the basis of computational feasibility, especially in the high dimensional case that we address.

2 The Long Memory Orthogonal Factor Conditional Volatility Model

We propose a model for volatility that embeds the univariate long memory EWMA model of Zumbach (2006) in an orthogonal factor structure. The long memory orthogonal factor (LMOF) conditional volatility model is developed in the spirit of the Orthogonal GARCH model of Alexander (2001) and the Generalised Orthogonal GARCH model of Van der Weide (2002). However, our model adopts a richer specification than these models. While the Orthogonal GARCH models allow for conditional volatility in the factors only, we assume that the volatilities of both the factors and the idiosyncratic shocks are time-varying and exhibit long memory behaviour. For both the factors and the idiosyncratic shocks, volatility is modelled using the univariate long memory LM-EWMA model. A significant advantage of the LM-EWMA model over other long memory volatility models is that it is much less computationally burdensome, which greatly facilitates its implementation in practice.

2.1 The Factor Specification

Consider a vector of returns on n assets $\mathbf{r}_t = (r_{1t}, r_{2t}, \dots, r_{nt})'$. In the factor framework, asset returns are linearly decomposed into a component that is correlated with a set of market-wide risk factors and an asset-specific component:

$$\mathbf{r}_t = \boldsymbol{\alpha} + \mathbf{B}\mathbf{f}_t + \boldsymbol{\varepsilon}_t \quad (1)$$

where $\mathbf{f}_t$ is a vector of k common risk factors with $k \ll n$, $\mathbf{B}$ is an $n \times k$ matrix of factor loadings, and $\boldsymbol{\varepsilon}_t$ is a vector of asset-idiosyncratic returns. The idiosyncratic shocks are uncorrelated with the factors $\text{cov}(\mathbf{f}, \boldsymbol{\varepsilon}) = \mathbf{0}$, a standard assumption in factor models, while they are allowed to be mildly cross-sectionally correlated with each other. The allowance for some correlation in the idiosyncratic components yields an approximate factor structure, first introduced by Chamberlain and Rothschild (1983). The conditional covariance matrix $\mathbf{H}_t$ can thus be represented as the combination of a common factor component and an idiosyncratic component:

$$\mathbf{H}_t = \mathbf{B}\boldsymbol{\Omega}_t\mathbf{B}' + \boldsymbol{\Sigma}_t \quad (2)$$

where $\boldsymbol{\Omega}_t$ is the covariance matrix of the common factors $\mathbf{f}_t$, and $\boldsymbol{\Sigma}_t$ is the covariance matrix of the idiosyncratic shocks $\boldsymbol{\varepsilon}_t$.

The literature shows that Principal Components Analysis (PCA) provides a consistent estimator of the space spanned by the ‘true’ latent factors (see, for example, Stock and Watson, 2002, Forni et al., 2000, Forni et al., 2005). While there is no reason to assume that the true factors are orthogonal, any arbitrary rotation of the factors that spans the original space is permissible, and PCA selects factors that are orthogonal. This obviates the need to model the off-diagonal elements of the covariance matrix $\mathbf{\Omega}_t$. We thus apply PCA to the panel of returns to estimate the factor loadings $\mathbf{B}$ and obtain the factors $\mathbf{f}_t$ and the idiosyncratic shocks $\mathbf{\varepsilon}_t$. For simplicity, we assume that the number of factors is constant. The factor covariance matrix is then a diagonal matrix with the variance of $f_{i,t}$ on the i^{th} diagonal: $\mathbf{\Omega}_t = \text{diag}\{\sigma_{f_{i,t}}^2\}$. As for the idiosyncratic components $\mathbf{\varepsilon}_t$, since they are only weakly cross-correlated, component-wise residuals can be obtained without much loss of information via univariate fitting: $\varepsilon_t \sim N(0, \Sigma_t)$ where Σ_t is an $n \times n$ diagonal matrix containing the conditional variances of each idiosyncratic series. Also, since our purpose is to forecast the covariance matrix, accounting for the weak correlation of the idiosyncratic shocks may not be useful (Luciani, 2014).

2.2 Incorporating Long Memory Volatility in the Factor Structure

In the LMOF model, the volatility of both the factors and the idiosyncratic shocks is assumed to be stationary, but to display long memory. We model the volatility of both components using the long memory LM-EWMA model of Zumbach (2006). Like the short memory EWMA model of JP Morgan (1994) on which it is based, the LM-EWMA model has a highly parsimonious specification, which greatly facilitates its implementation. In the LM-EWMA model, conditional volatility is defined as the weighted average of K standard (i.e. short memory) EWMA processes over increasing time horizons:

$$\sigma_t^2 = \sum_{k=1}^K w_k \sigma_{k,t}^2, \quad (3)$$

where $\sigma_{k,t}^2$ is a EWMA process given by

$$\sigma_{k,t}^2 = \mu_k \sigma_{k,t-1}^2 + (1 - \mu_k) r_{t-1}^2. \quad (4)$$

The decay factor μ_k is characterised by a geometric time horizon τ_k such that $\mu_k = \exp(-1/\tau_k)$, with $\tau_k = \tau_1(\rho)^{k-1}$ for $k = (1, \dots, K)$. Zumbach (2006) sets the value of ρ to $\sqrt{2}$.

The long memory of the volatility process in Equation (3) is determined by the weights w_k , which are assumed to decay logarithmically:

$$w_k = \frac{1}{C} \left(1 - \frac{\ln(\tau_k)}{\ln(\tau_0)} \right), \quad (5)$$

with the normalization constant C such that $\sum_k w_k = 1$. The conditional long memory volatility is therefore parsimoniously parameterised by just three factors: τ_1 (the shortest time scale at which the volatility is measured, i.e. the lower cut-off), τ_K (the upper cut-off, which increases exponentially with the number of components K), and τ_0 (the logarithmic decay factor). Zumbach (2006) suggests parameter values of $\tau_0 = 1560$ days = 6 years, $\tau_1 = 4$ days and $\tau_K = 512$ days, which is equivalent to $K = 15$.

Since the long memory EWMA volatility is the sum of EWMA processes, volatility forecasts are straightforward to obtain using a recursive procedure. The EWMA volatility process in Equation (4) can be expressed in the exponentially weighted form:

$$\sigma_{k,t}^2 = (1 - \mu_k) \sum_{i=1}^{\infty} \mu_k^i r_{t-i}^2. \quad (6)$$

Substituting Equation (6) into Equation (3), the long memory EWMA volatility process can be written as:

$$\sigma_t^2 = \sum_{k=1}^K w_k (1 - \mu_k) \sum_{i=1}^{\infty} \mu_k^i r_{t-i}^2 = \sum_{i=1}^{\infty} \lambda(i) r_{t-i}^2 \quad (7)$$

with the logarithmically decaying weights $\lambda(i) = \sum_k w_k (1 - \mu_k) \mu_k^i$ and $\sum_i \lambda(i) = 1$. Under the assumption of serially uncorrelated returns, the h -step-ahead cumulative volatility forecast, given the information set F_t at time t , is equal to:

$$\sigma_{t+1:t+h}^2 = \sum_{i=0}^T \lambda(h, i) r_{t-i}^2 \quad (8)$$

with the weight $\lambda(h, i)$ given by

$$\lambda(h, i) = \sum_{k=1}^K \sum_{j=1}^h w_{j,k} \frac{(1 - \mu_k)}{1 - \mu_k^T} \mu_k^i \quad (9)$$

where T is the cut-off time³, $w_{j,k}$ is the k^{th} element of vector $\mathbf{w}_j = \mathbf{w}'(\mathbf{M} + (\mathbf{I} - \boldsymbol{\mu})\mathbf{w}')^j$, $\mathbf{w} = (w_1, w_2, \dots, w_K)'$, $\boldsymbol{\mu}$ is the vector of μ_k , $\mathbf{M}$ is the diagonal matrix consisting of μ_k , and $\mathbf{I}$ is the unit vector. Since $\sum_k w_k = 1$, we obtain $\sum \lambda(h, i) = 1$. Note that when $K = 1$, we have $w = 1$ and the long memory EWMA forecast function collapses to a standard short memory EWMA forecast function with forecast weight $\lambda(h, i) = h(1 - \mu_k)\mu_k^i / (1 - \mu_k^T)$. Since the weight $\lambda(h, i)$ is independent of the data, the forecast in Equation (8) is straightforward to compute. As with the standard EWMA model, the univariate long memory LM-EWMA model can be easily extended to the multivariate setting (see Zumbach, 2011, Harris and Nguyen, 2013).

In the LMOF model, the univariate LM-EWMA model is used to estimate the volatilities of the factors $\sigma_{f,t+1}^2$ and the idiosyncratic shocks $\sigma_{\varepsilon,t+1}^2$, from which the covariance matrices of the factors $\Omega_t = \text{diag}\{\sigma_{f_{id}}^2\}$ and of the shocks $\Sigma_t = \text{diag}(\sigma_{\varepsilon,t}^2)$ are easily constructed. Ω_t and Σ_t are then combined to estimate the conditional covariance matrix $\mathbf{H}_t$. Under the assumption of serially uncorrelated factors and residuals, the h -step ahead forecast of the conditional covariance matrix of returns is given by

$$\mathbf{H}_{t+1:t+h} = \mathbf{B}\Omega_{t+1:t+h}\mathbf{B}' + \Sigma_{t+1:t+h} \quad (10)$$

with $\Omega_{t+1:t+h} = \text{diag}\{\sigma_{f_{id,t+1:t+h}}^2\}$ and $\Sigma_{t+1:t+h} = \text{diag}(\sigma_{\varepsilon,t+1:t+h}^2)$.

3 The Economic Value of Conditional Volatility Models

We examine the economic performance of the LMOF model using the dynamic volatility timing framework of Fleming et al. (2001). The LMOF model is evaluated against two long memory models (the fully estimated multivariate LM-EWMA model and the CGARCH-DCC model) and two short memory models (the standard EWMA model and the GARCH-DCC model).

3.1 The Dynamic Volatility Timing Framework

³ Zumbach (2011) suggests that for many practical applications, the memory length T is of the order of one to two years. In the empirical application, we set T equal to the estimation window length.

Suppose that an investor allocates fraction $\mathbf{w}_t$ of his wealth to the n risky assets and the remainder $(1 - \mathbf{w}_t' \mathbf{1})$ to the risk-free asset, where $\mathbf{1}$ is the $(n \times 1)$ unit vector. In the mean-variance optimization framework, the investor maximizes his expected utility U_{t+1} :

$$\max_{\mathbf{w}_t} \left\{ E(U_{t+1}) = \mu_{p,t+1} - \frac{\lambda}{2} \sigma_{p,t+1}^2 \right\} \quad (11)$$

where $\mu_{p,t+1} = \mathbf{w}_t' \boldsymbol{\mu}_{t+1} + (1 - \mathbf{w}_t' \mathbf{1}) r^f$ is the portfolio's expected returns, $\sigma_{p,t+1}^2 = \mathbf{w}_t' \mathbf{H}_{t+1} \mathbf{w}_t$ is the portfolio's expected variance, $\boldsymbol{\mu}_{t+1}$ is the vector of expected returns, $\mathbf{H}_{t+1}$ is the expected conditional covariance matrix of returns, r^f is the risk-free rate and λ is the risk aversion coefficient. Following Fleming et al. (2001), expected returns are assumed constant ($\boldsymbol{\mu}_{t+1} \equiv \boldsymbol{\mu}$) so that we can specifically examine the economic value of volatility predictability. In the empirical study, we assume a risk free rate of 4% and a risk aversion coefficient of 1. We consider alternative values of λ in the robustness test. Assuming short sales are allowed, and that there are no transaction costs, the solution to this optimization problem is:

$$\mathbf{w}_t^* = \frac{1}{\lambda} \mathbf{H}_{t+1}^{-1} (\boldsymbol{\mu} - r^f). \quad (12)$$

If the covariance matrix is constant, the optimal weights will be constant over time, which is the static strategy. However, if the investor believes that the covariance matrix is time-varying, he will follow a dynamic strategy by changing the optimal weights based on his forecast of the conditional covariance matrix. By comparing the performance of the static and dynamic strategies, we can evaluate the economic value of volatility timing. We consider dynamic strategies based on three long memory volatility models (the LMOF, multivariate LM-EWMA and CGARCH(1,1)-DCC models) and two short memory volatility models (the multivariate EWMA and GARCH(1,1)-DCC models).

3.2 Performance Evaluation

We evaluate the performance of the dynamic portfolio strategies using three common measures. First, the out-of-sample Sharpe ratio of each strategy is calculated as the sample mean of the realised portfolio excess returns over the risk free rate divided by their sample standard deviation, $SR = (\mu_p - r^f) / \sigma_p$. While the Sharpe ratio is the most commonly used performance measure, Han (2006) argues that it may be misleading in the sense that it does not take into account the time-varying conditional volatility. Using the realised sample standard

deviations, the ex post Sharpe ratio may overestimate the conditional risk that an investor faces in a dynamic strategy, and hence underestimate the strategy performance. Therefore, we additionally consider a related measure that compares the relative performance of the dynamic and static portfolios on an equal basis in terms of risk. This is the $M2$ measure of Modigliani and Modigliani (1997), which is defined as the excess return that the dynamic strategy would earn if it had the same risk as the static benchmark:

$$M2 = \frac{\sigma_b}{\sigma_p} (\mu_p - r^f) - (\mu_b - r^f) \quad (13)$$

where μ_s, σ_s and μ_d, σ_d are the out-of-sample means and standard deviations of the static and dynamic portfolios, respectively. Note that this measure is directly related to the Sharpe ratios of the two strategies, since $M2 = \sigma_s (SR_d - SR_s)$.

The third measure is the performance fee suggested by Fleming et al. (2001) and now widely used in the literature on measuring the economic value of investment. Consider the quadratic utility function:

$$U(W_{t+1}) = W_{t+1} - \frac{a}{2} W_{t+1}^2, \quad (14)$$

where W_{t+1} is the investor's expected wealth. Assume that each day the initial wealth of the investor is fixed at W_0 . Assume also that the investor's coefficient of relative risk aversion,

$\gamma_t = \frac{aW_t}{1-aW_t}$, is constant. The investor's average utility function is given by

$$\bar{U}(\cdot) = W_0 \left(\sum_{t=0}^{T-1} R_{p,t+1} - \frac{\lambda}{2(1+\gamma)} R_{p,t+1}^2 \right). \quad (15)$$

By setting γ to be constant, Equation (15) represents a second order approximation of a non-quadratic utility function. Constant relative risk aversion also implies that expected utility is linearly homogenous in wealth. The performance fee Δ is defined as the maximum fee that the investor would be willing to pay to switch from the static strategy to the dynamic strategy, without being worse off in terms of utility. To estimate this fee, we find the value of Δ that equates the realised average utilities for the two alternative portfolios:

$$\sum_{t=0}^{T-1} (R_{d,t+1} - \Delta) - \frac{\gamma}{2(1+\gamma)} (R_{d,t+1} - \Delta)^2 = \sum_{t=0}^{T-1} R_{s,t+1} - \frac{\gamma}{2(1+\gamma)} R_{s,t+1}^2, \quad (16)$$

where $R_{d,t+1}$ and $R_{s,t+1}$ are the gross realised returns of the dynamic and static strategies, respectively. In the empirical analysis, we report the annualised performance fees in basis points for values of the constant relative risk aversion parameter $\gamma = 1$ and $\gamma = 5$.

3.3 Transaction costs

A volatility timing strategy involves the regular updating of the positions in individual assets within the portfolio, which incurs transaction costs that will tend to offset the gains that arise from the dynamic strategy. Following Han (2006), we estimate the breakeven transaction cost τ^{be} , defined as the transaction cost that makes the investor indifferent between the dynamic and the static strategies in term of utility. If the investor's actual transaction costs are lower than the breakeven transaction cost, he will be better off with the dynamic strategy; otherwise he should follow the static benchmark. Han sets the transaction cost equal to a fixed percentage (τ) of the value traded for all stocks. The costs for the static and dynamic strategies are then given by

$$\tau \left| w - \frac{w(1 + r_{p,t+1})}{w(r_{p,t+1} - r^f) + r^f + 1} \right| \quad (17)$$

$$\text{and } \tau \left| w_t - \frac{w_t(1 + r_{p,t+1})}{w_t(r_{p,t+1} - r^f) + r^f + 1} \right|, \text{ respectively.} \quad (18)$$

The breakeven transaction cost is computed by equating the utilities of the static and dynamic strategies as in Equation (16) after taking account of the costs. The higher the breakeven transaction cost, the more easily the dynamic trading strategies can be implemented. Since the breakeven transaction cost is a proportional cost paid every time the portfolios are rebalanced, we report this cost in basis points at the rebalancing frequency, e.g., for a daily rebalanced portfolio, we report the cost in daily basis points. The breakeven transaction cost is only estimated when the performance fee in Equation (16) is positive.

4 Empirical Methodology

4.1 Data Description

We construct an international stock and bond dataset, comprising 21 stock indices from the FTSE All-World indices and 13 five-year average maturity bond indices. The 21 stock indices and 13 bond indices include all of the major global stock and government bond markets. The same dataset is employed by Engle and Colacito (2006) and Harris and Nguyen (2013). All of

the data are from Datastream and converted to US dollar denominated prices. Following Engle and Colacito (2006), we use weekly returns to avoid the problem of non-synchronous trading. Weekly returns are calculated as the log price difference using Friday-to-Friday closing prices. The dataset covers the period from 1 January 1988 to 31 December 2013, yielding a total of 1355 observations. Table 1 reports the descriptive statistics for the international series. For all countries for which both stock and bond indices are present, the stock index has a higher return and higher risk than the corresponding bond index. Returns are leptokurtic and, in most cases, negatively skewed. The international stock markets are relatively highly correlated, as are the international bond markets. The average correlation coefficient among the 21 stock market return series is 0.57, while among the bond market return series it is 0.59. However, the stock and bond markets as a whole have an average correlation coefficient of only 0.22.

[Insert Table 1 here]

Tests are conducted to confirm the evidence of long memory dynamics in the volatility of the stock and bond index returns, the results of which are also reported in Table 1. The parametric FIGARCH model is estimated for the whole sample, suggesting the presence of long memory volatility in all series. We also apply the semi-parametric long memory GPH tests of Geweke and Porter-Hudak (1983) for both squared and absolute return series. To estimate the GPH operator, we use the recommended bandwidth m equal to the square root of the sample size ($m = 77$). The GPH test generally confirms the presence of long memory in the volatility of the stock and bond return series. The results also suggest that stock return volatility has longer memory than bond return volatility, and that the memory of absolute returns is consistently longer than that of squared returns, a feature first identified by Taylor (1986).

4.2 Determining the Number of Factors

We assume that the number of factors is constant for the whole sample and employ three different methods to determine the number of factors. First, we estimate the percentage of variance explained by the eigenvalues of both the variance-covariance matrix and the spectral density matrix, the results of which are reported in Table 2. Second, we adopt the criterion of Alessi et al. (2010). This is a generalisation of the information criterion of Bai and Ng (2002), which selects the number of factors by minimizing the variance explained by the idiosyncratic component, while imposing a penalty function to avoid over-parameterization. Alessi et al. (2010) modify the Bai and Ng criterion by adding a multiplicative tuning parameter, c , in the penalty function, based on the procedure proposed by Hallin and Liska (2007) in the context

of dynamic factor models. Based on random subsamples, they obtain the behaviour of the information criterion for different values of c . From that they identify the smallest values of c for which the criterion is a constant function of the subsamples. As the implementation of the Alessi et al. criterion may be sensitive to the selection of subsamples and maximum input factors (k_{max}), we run the criterion 100 times for each $k_{max} = \overline{5,10}$. The criterion identifies four common factors in 417 out of 600 trials (69.5%) and two common factors in 183 out of 613 trials (30.5%). Finally, we employ the Onatski (2010) test, which is based on the fact that the ‘systematic’ eigenvalues (the common factors) diverge to infinity while the ‘idiosyncratic’ eigenvalues cluster around a single point. The estimator separates the diverging eigenvalues from the cluster and counts their number. Onatski argues that their estimator works well in small samples, which is the main advantage over to the Bai and Ng estimators. We apply their estimator on the grid $k_{max} = \overline{5,20}$. The estimator identifies four common factors for all $r_{max} > 8$, and two factors for $k_{max} = \overline{5,7}$.

For most of the trials, the Alessi et al. criterion and the Onatski test identify four common factors, which explain around 72.1% of the total variance in the variance-covariance matrix and 75.6% of the total variance in the spectral density matrix. In other trials, they choose two common factors. However, to evaluate the sensitivity of the choice of the number of factors, we employ two, four and five common factors in the empirical analysis.

[Insert Table 2 here]

4.3 The Estimation Process

We compare the economic value of the LMOF model against that of a wide range of multivariate conditional volatility EWMA and GARCH models, using the volatility timing framework of Fleming et al. (2001). The benchmark models include two multivariate long memory models – the multivariate LM-EWMA of Zumbach (2011) and the component CGARCH model of Engle and Lee (1999) implemented using the DCC framework – and two widely used short memory volatility models – the Riskmetrics’ EWMA model of JPMorgan (1994) and the GARCH-DCC model.

The whole sample is divided into an initialisation period and a forecast period. The initialisation period is from 1 Jan 1988 to 31 Dec 1993 (312 weekly observations) and the forecast period from 1 Jan 1994 to 31 Dec 2013 (1043 observations). Expected returns are assumed to be constant and are set equal to the sample means estimated using the initialisation period. The investor rebalances his portfolio periodically, based on changes in the estimated conditional

covariance matrix. We consider both weekly and monthly rebalancing. The conditional volatility model is first estimated using the initialisation period, and used to generate a one step ahead forecast of the conditional covariance matrix. This is used to compute the optimal portfolio weights for the following period, with which we calculate the realised portfolio return for that period. The estimation window is then rolled forward one period, the model re-estimated, a new forecast of the conditional covariance matrix is made and the portfolio is rebalanced, and so on until the end of the sample is reached. The realised performance of each dynamic portfolio constructed with the different conditional volatility models is compared with that of the ex ante optimal static portfolio, based on the sample mean and covariance matrix estimated over the initialisation period. We also compare the dynamic portfolios against the equally weighted portfolio, which is a widely used benchmark in practice.

5 Empirical Results

5.1 Performance Analysis of the Dynamic Asset Allocation Strategies

Table 3 reports the out-of-sample performance of the international stock and bond portfolios with weekly and monthly rebalancing. We report results for the dynamic portfolios computed using three long memory models (the LMOF model, the LM-EWMA model and the Component CGARCH(1,1)-DCC model), and the two short memory models (the EWMA model and the GARCH(1,1)-DCC model). We also report results for the two benchmark portfolios (the static portfolio and the 1/N portfolio). The ‘*LMOF k* ’ portfolio refers to the portfolio constructed using the LMOF model with k factors. It is clear that all the dynamic portfolios significantly outperform both the static and equally weighted portfolios. The conditional volatility models consistently produce portfolios with higher Sharpe ratios and $M2$ measures, and positive excess returns. The passive investor would be willing to pay an annualised performance fee Δ_1 of at least 132 *bps* to switch from the static to the dynamic strategies. The findings are consistent with the existing literature (see, for example, Fleming et al., 2001, 2003, Han, 2006), confirming the value of volatility timing in asset allocation.

Among the dynamic portfolios, those that are based on the LMOF model have higher Sharpe ratios and $M2$ measures, than those constructed with the other multivariate volatility models, both short and long memory, and at both investment horizons. The results are robust to the number of common factors, although among the three LMOF portfolios, the LMOF4 portfolio dominates. The LMOF portfolios also have higher performance fees than those based on the GARCH-DCC and LM-GARCH-DCC models, but lower than those based on the EWMA and

LM-EWMA models. For example, with the relative risk aversion of $\gamma = 1$, the performance fee for the weekly rebalanced LMOF4 portfolio is 336 *bps*, much lower than the 861 *bps* for the corresponding LM-EWMA portfolio. It is interesting that the simple EWMA and LM-EWMA models outperform more sophisticated models such as the GARCH-DCC and CGARCH-DCC models.

Once we consider transaction costs, however, the LMOF models dominate all of the other models, including the EWMA and LM-EWMA models. For example, a weekly trader with $\gamma = 1$ is only better off with the LM-EWMA portfolio if his realised transaction cost is lower than 5 *bps*, compared to that of 28 *bps* if he employs the LMOF4 portfolio. In terms of transaction costs, only the LMOF model generates portfolios that are feasible in practice. The breakeven transaction costs of a monthly trader are much higher than those of a weekly trader due to less frequent rebalancing.

[Insert Table 3 here]

5.2 Estimation Error in Expected Returns

Fleming et al. (2001) suggest that using a single vector of expected returns may be inappropriate. We hence follow their recommendation and consider a range of expected returns that are generated via a block bootstrap procedure. We generate an artificial sample of 4,000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. We then calculate the unconditional mean and covariance matrix of this artificial return series. The unconditional mean from the bootstrap, together with our forecasts of the conditional covariance matrix, are used to estimate the weights of the optimal dynamic portfolio. We apply these weights to the actual returns to calculate the realised portfolio return in the following period. The benchmark static portfolio is constructed using the unconditional mean and covariance matrix from the bootstrap. We repeat this procedure with 1,000 trials, studying the economic gains from volatility timing across a wide range of plausible vectors of expected returns.

Table 4 reports the average results across the 1,000 bootstrap vectors of expected returns for the international stock and bond dataset. Consistent with the results obtained using the single expected return vector, the investor is generally better off switching from the static strategies to the dynamic strategies, and among the dynamic strategies, to the LMOF model. The LMOF model generates portfolios with positive average excess returns and performance fees, and higher Sharpe ratios than the static unconditional covariance matrix estimator in all trials. The

LMOF model also dominate other dynamic models in terms of the Sharpe ratio and average excess return. However, it fails to outperform the EWMA and LM-EWMA models in terms of performance fees owing to the lower realised portfolio returns. Among the factor models, the LMOF4 model again performs best. For example, the LMOF4 model generates weekly rebalanced portfolios with a Sharpe ratio of 1.269, compared to 1.141 for the LMOF5 model. The LMOF model may again be the only feasible strategies in practice with their high breakeven transaction costs. The results also suggest that the simpler EWMA and LM-EWMA models again outperform the more sophisticated GARCH-DCC and CGARCH-DCC models, which underperform even the two static strategies.

[Insert Table 4 here]

5.3 Sensitivity to the Risk Aversion Coefficient

The results reported in the previous section are based on a risk aversion coefficient of $\lambda = 1$. In this section, we evaluate the performance of the dynamic strategies for alternative values of the risk aversion coefficient from 1 to 5. In each case, we repeat the experiment with 1,000 bootstrap vectors of expected returns. Table 5 compares the static and the dynamic strategies across the different risk aversion coefficients for the international stock and bond portfolios. To save space, we report the results only for the LMOF4 model and the LM-EWMA model. Expectedly, when the investor is more risk averse (i.e. for higher values of λ), he will choose more conservative portfolios with lower risk and lower expected return. The Sharpe ratios are approximately the same across all risk aversion levels, with a slight difference arising from the bootstrap procedure. Consistent with the previous results, the dynamic portfolios (and especially the factor portfolio) generate higher Sharpe ratios than the static portfolios for most bootstrap vectors. The average Sharpe ratio of the factor portfolio is as much as three times as high as that of the static portfolio. The investor would be willing to pay an annualised performance fee of around 50 to 200 *bps* to switch from the static portfolios to the LMOF4 portfolios. Compared to the LM-EWMA model, the LMOF4 model produces portfolios with higher Sharpe ratios, lower performance fees due to lower realised portfolio returns, but higher breakeven transaction costs as a result of less rebalancing.

[Insert Table 5 here]

5.4 Sub-period Performance

Sub-period performance is of interest as the performance of the dynamic strategies is likely to vary according to market conditions. Table 6 reports the out-of-sample performance of the

static and dynamic strategies over different years. The bootstrap procedure is again used to control for estimation error in expected returns. To save space, we report results only for the static portfolio and the dynamic portfolio based on the LMOF4 model, with daily rebalancing. Unsurprisingly, the optimal portfolios, both static and dynamic, closely track the state of the market. In particular, the realised performance of the two strategies deteriorates significantly in 1995 (the Mexican crisis), 1998 (the Asian crisis), 2001-2002 (the dotcom bust), 2008-2009 (the global crisis) and 2011 (the European debt crisis).

The results suggest that the dynamic portfolio outperforms the static portfolio in 16 out of 20 years. It is also interesting to note that the dynamic strategy performs reasonably well in market downturns. For example, the realised returns of the LMOF4 portfolio increased from 7.55% in 2007 to 7.87% in 2008, in sharp contrast to the big drop from 4.28% to -3.65% of the static portfolio. In terms of performance fees, an investor would be willing to pay more to switch to the dynamic strategy in market downturns than in normal market conditions. This is because the conditional volatility model allows investors to better estimate the high volatility that is associated with market downturns, thus allowing them to exploit the benefits of volatility timing during these periods. In normal market conditions, however, the results are mixed, and there are some years in which the investor will be better off staying with the static strategy.

[Insert Table 6 here]

5.5 Application to the Dow Jones Industrial Average Dataset

We additionally consider a higher frequency, high dimensional dataset, comprising the components of the Dow Jones Industrial Average (DJIA) index as of 31 December 2013. Daily data are collected from the Centre for Research in Security Prices from 1 June 1999 to 31 December 2013. We exclude Visa, which went public only in March 2008. Returns are calculated as the log price difference over consecutive days. All days on which the market was closed are excluded from the sample, yielding 3671 observations. A similar procedure is applied to determine the number of factors. The Bai and Ng (2002) criterion and Alessi et al. (2010) modified criterion both identify two common factors, while the Onatski (2010) estimator identifies only one common factors. In the empirical study, we report the results for both one and two common factors.

The whole sample is again divided into an initialisation period (1 June 1999 to 31 May 2005, 1510 daily observations) and a forecast period (1 June 2005 to 31 Dec 2013, 2161 observations). The bootstrap procedure is employed to account for estimation error in expected

returns. We use all the multivariate volatility models to construct the dynamic covariance matrix. We additionally employ the FIGARCH(1,d,1)-DCC model of Baillie et al. (1996). This model is excluded in the international stock and bond experiment since its estimation requires a prohibitively high upper lag cut-off. Following standard practice in the literature, we set the truncation lag for the FIGARCH model equal to 1000.

The results reported in Table 7 show that the two LMOF models do not perform as well in the DJIA dataset as in the international stock and bond dataset. It is possible that the greater level of noise associated with the higher sampling frequency results in a lower degree of measurable comovement among the assets, and hence reduces the significance of the factor structure. However, the LMOF model still generally dominates the static strategy and the dynamic DCC family for all performance measures and rebalancing frequencies. For example, an investor with a constant relative risk coefficient of $\gamma = 1$, who rebalances his portfolio weekly, would be willing to pay an average performance fee of between 60 and 68 *bps* to switch from the static strategy to the LMOF strategy. Again, the parsimonious EWMA and LM-EWMA models perform very well. They generally outperform the more sophisticated GARCH-DCC and CGARCH-DCC models. The EWMA and LM-EWMA models even generate portfolios with higher performance fees than the LMOF model and outperforms it in terms of the Sharpe ratios with monthly rebalancing. The LMOF model, however, is more feasible in practice with their higher breakeven transaction costs that result from lower turnover. The high degree of parameterization of the FIGARCH model evidently hinders its performance and it underperforms all other models. We also evaluate the performance of the dynamic strategies using different values of the risk aversion coefficient and the results are very similar.

[Insert Table 7 here]

5.6 Comparison with Static Factor Benchmarks

We further evaluate the gains of allowing for long memory volatility in a factor structure by comparing the LMOF model with (a) the unconditional factor model and (b) the conditional short memory factor EWMA model. The unconditional *Factor* k model is the traditional rolling window factor model with k factors, estimated using Principle Component Analysis. The conditional short memory factor EWMA (SMOF) model is constructed in a similar approach to the LMOF model, but the factor and residual volatilities are estimated using the standard EWMA model rather than the LM-EWMA model. Again, the bootstrap procedure is employed to account for estimation error in expected returns. To save space, we only report the results

for four common factors with the international stock and bond dataset, and two common factors with the DJIA dataset in Table 8. In the international stock and bond dataset, the long memory factor models consistently produce portfolios that are superior to those produced by both the traditional unconditional factor models and the conditional short memory factor EWMA models. The LMOF portfolios consistently yield higher Sharpe ratios and abnormal returns in most of the trials. The investor is also willing to pay more to switch from the static to the LMOF portfolios than from the static to the unconditional factor portfolios. However, there is only a small difference between the long memory LMOF portfolios and the SMOF portfolios in terms of performance fees, where the SMOF portfolios perform marginally better. Allowing for long memory in the factor structure requires more rebalancing, leading to smaller breakeven transaction costs. These transaction costs are, nevertheless, still high enough make implementation of the LMOF portfolios viable in practice. The performance of the factor models is less impressive in the case of the DJIA dataset. Although the LMOF model still outperforms the unconditional factor model, there is no significant difference between portfolios generated by the LMOF and SMOF models. In terms of transaction costs, the two conditional volatility factor models yield lower breakeven transaction cost than the unconditional factor models, which may make them undesirable for daily trading. However, they are again high enough for the lower trading frequency of weekly and monthly rebalancing.

[Insert Table 8 here]

5.7 Comparison with a Generalised Dynamic Factor Benchmark

Recently, generalised dynamic factor models, in which the common component is dynamic and is allowed to have a moving average representation, have become a popular tool in finance and macroeconomics (see Forni et al., 2000, Forni et al., 2005, Hallin and Lippi, 2013, Barigozzi and Hallin, 2016, to quote a few). In this section, we compare the performance of our static LMOF factor model with the generalised dynamic factor model of Alessi et al. (2009).

Let $\mathbf{r}_t = (r_{1t}, r_{2t}, \dots, r_{nt})'$ is the vector of asset returns. In the Generalised Dynamic Factor Model (GDFM) introduced by Forni et al. (2000), each return process is the sum of a common component χ_{it} and an idiosyncratic component ξ_{it} . The common component is driven by q dynamic common shocks which are loaded with potentially different coefficients and lags:

$$\mathbf{r}_t = \boldsymbol{\chi}_t + \boldsymbol{\xi}_t = \mathbf{B}(L)\mathbf{u}_t + \boldsymbol{\xi}_t \quad (19)$$

where $\mathbf{u}_t$ is a q -dimensional orthonormal white noise vector, $\mathbf{B}(L)$ is a one-sided $n \times q$ absolutely summable matrix polynomial. With some transformation, Equation (19) can be rewritten in the static form:

$$\mathbf{r}_t = \Lambda \mathbf{f}_t + \xi_t \quad (20)$$

where $\mathbf{f}_t$ is an r -dimensional *static* factor with $\mathbf{D}(L)\mathbf{f}_t = \mathbf{u}_t$.

Alessi et al. (2009) allow the dynamic factors to be conditional heteroscedastic, $\mathbf{u}_t \sim N(0, \mathbf{Q}_t)$ with $\mathbf{Q}_t$ being modelled with the BEKK or DCC structure, while the idiosyncratic component follows a univariate ARMA-GARCH process. This yields the Dynamic Factor GARCH (DFGARCH) model.

We employ the Hallin and Liska (2007) criterion to determine the number of dynamic factors and it identifies one common dynamic factor for both the international stock and bond, and the DJIA datasets.⁴ The dynamic factor may represent the world factor for the international stock and bond dataset, and the US market factor for the DJIA dataset. This may be a general feature of financial data when many other studies have also obtained one dynamic factor (Luciani and Veredas, 2015, Barigozzi and Hallin, 2015). Forni et al. (2005) develop a static presentation of their dynamic factor model in which lagged dynamic factors are introduced as additional static factors $r = q(s + 1)$ where s is the lag of the dynamic factors (see Equation 20), thus it normally requires a higher number of static factors than that of dynamic factors to explain the same percentage of variances.

The GARCH(1,1)-DCC structure is applied to estimate the covariance matrix for the dynamic factor in the Alessi et al. (2009) DFGARCH model. Following their suggestion, we implement the Kalman filter to improve forecast accuracy. We employ two DFGARCH models. Following Alessi et al., we estimate the first DFGARCH model with the static number of factors used in the LMOF model and the number of dynamic factors identified by the Hallin and Liska criterion. We additionally consider a second DFGARCH model in which we set the number of dynamic factors equal to that of the static factors. We again employ the bootstrap procedure to estimate vectors of expected returns. Table 9 compares the performance of the static portfolios and the dynamic portfolios constructed with the two DFGARCH models and the LMOF models. The DFGARCH(s, d) model is the Dynamic Factor GARCH model with s static factors

⁴ Employing the FTSE100 dataset, Alessi et al. (2009) identify 5 static factors and only 2 dynamic factors.

and d dynamic factors. Interestingly, the dynamic DFGARCH model does not dominate the LMOF models in these experiments.⁵ With the international stock and bond dataset, the DFGARCH model perform quite well in terms of performance fees due to high realised returns. However, it performs much worse in terms of the Sharpe and M2 ratios. It is interesting to see that the second DFGARCH(4,4) model outperforms the first DFGARCH(4,1) model. A similar result applies to the DJIA dataset. Allowing for more dynamic factors may capture more comovement in stock returns, yielding better forecast performance. The forecast power of the DFGARCH model drops significantly with the DJIA dataset, where the DFGARCH portfolios cannot even beat the static portfolios. The DFGARCH model also underperforms the GARCH(1,1)-DCC and the traditional static factor portfolios. This is at odds with the perception that a fully specified dynamic factor model should yield better forecasts. As argued by Boivin and Ng (2005), the forecasting equation need not be a full-blown factor model and hence the stronger the adherence to a factor structure, the more likely that the estimation errors of the factor estimates will enter the forecasts. Static factor models may be more favourable due to the minimal factor structure imposed on the forecasting equation and the simplicity in its implementation. The parsimony of static factor models may offset the possible gains from more correctly specified, yet more complex dynamic factor models. This issue deserves attention in future research.

[Insert Table 9 here]

5.8 Other Robustness Tests

Sensitivity to the Estimation Windows

As the factor loadings estimated using Principal Component Analysis may be sensitive to the sample length used in their estimation, we investigate the performance of the strategies using a range of estimation windows. In particular, we consider 4, 6, 8 and 10 years of weekly data for the international stock and bond dataset, and 2, 4, 6 and 8 years of daily data for the DJIA dataset. The analysis is again conducted with the bootstrap vectors of expected returns. Figure 2 shows the average Sharpe ratios of the LMOF, LM-EWMA, SMOF and unconditional factor models using different estimation windows, for the two portfolios with weekly and monthly

⁵ We also employed the two-step General Dynamic Factor (GDFM) model of Barigozzi and Hallin (2015) as another benchmark. They propose a two-step procedure that first decomposes returns into ‘level-common’ and ‘level-idiosyncratic’ components, and then decomposes volatilities into four different components: common and idiosyncratic components of level-common innovations and common and idiosyncratic components of level-idiosyncratic innovations. For forecasting purposes, the univariate GARCH model is employed to forecast the four volatility components, from which the covariance matrix is easily obtained. The results also suggest that the two-step GDFM model fails to outperform the static factor LMOF model.

rebalancing. The LMOF portfolios generally dominate the LM-EWMA portfolios across all estimation windows. The LMOF model is also found to outperform the unconditional factor model. While the LMOF model performs better than the SMOF model in the international stock and bond dataset, there is only a small difference in the performance of the two models in the DJIA dataset. It is notable that the Sharpe ratios of the factor models tend to decline with the estimation window length, suggesting that data from the distant past is less relevant for estimation of the factor loadings.

[Insert Figure 1 here]

Time-Varying Number of Factors

So far, we assume a constant number of factors for the whole sample. However, the systematic risk dimension of financial assets is likely to be time-varying since it depends on economic regimes. We now relax this assumption and allow for a time-varying number of factors. In each rolling period, we estimate the number of factors using the Alessi et al. (2010) criterion, and input this number into the LMOF model. However, the results suggest that the LMOF portfolios with a time-varying number of factors fail to outperform those in which the number of factors is constant. Figure 2 compares the two LMOF portfolios. The figure plots the realised Sharpe ratios for 1,000 trials of the bootstrap experiment, with each dot representing a separate trial. The distribution of the Sharpe ratio is clearly below the 45-degree line, suggesting the outperformance of the LMOF portfolios with a constant number of factors. An experiment with the DJIA dataset, although not reported here, produces similar results. We additionally employ the Onatski (2010) test to determine the number of factors and the findings are similar. The LMOF model with a time-varying number of factors is still consistently superior to other multivariate conditional volatility models across all investment horizons. The time-varying number of factors may better capture the risk dimension of financial returns. However, the estimation error in estimating the time-varying number of factors may outweigh their benefits. Since the Alessi et al. (2010) and Onatski (2010) estimators may be sensitive to r_{max} , and also to the choice of subsamples, running these estimators once in each period may not be ideal. However, it would be too computationally burdensome to apply the procedure detailed in Section 4.2 for each period. It would be interesting to investigate this issue in greater detail.

[Insert Figure 2 here]

6 Conclusion

In this paper, we develop a dynamic factor long memory conditional volatility model (the LMOF model) that incorporates the long memory volatility behaviour in an orthogonal factor structure. The new model captures the high persistence of financial volatility that is observed in practice, while reducing the estimation error that arises from modelling high dimensional covariance matrices. We evaluate the economic benefits of the new model in the volatility timing framework. Dynamic portfolios based on the LMOF model are compared against a wide range of benchmark portfolios, including the static unconditional portfolio and other dynamic short memory and long memory portfolios. Consistent with the literature, we show that investors should be willing to pay to switch from static strategies to dynamic volatility timing strategies. Among the dynamic portfolios, the LMOF volatility model generally generate superior portfolios to other short and long memory volatility models. Our results also suggest that combining long memory volatility and the factor structure yields better results than employing either long memory volatility or factor structure alone. These results apply to both datasets, and are robust to estimation error in expected returns, the choice of risk aversion coefficient, the length of estimation window, and sub-period performance.

References

- ALESSI, L., BARIGOZZI, M. & CAPASSO, M. 2009. Estimation and forecasting in large datasets with conditionally heteroskedastic dynamic common factors. *ECB Working Paper Series*, 1115.
- ALESSI, L., BARIGOZZI, M. & CAPASSO, M. 2010. Improved penalization for determining the number of factors in approximate factor models. *Statistics & Probability Letters*, 80, 1806-1813.
- ALEXANDER, C. 2001. Orthogonal Garch. *Mastering Risk*, 2, 21--38.
- AMENC, N. & MARTELLINI, L. 2002. Portfolio Optimization and Hedge Fund Style Allocation Decisions. *The Journal of Alternative Investments*, 5, 7-20.
- ANDERSEN, T. G., BOLLERSLEV, T., DIEBOLD, F. X. & EBENS, H. 2001. The distribution of realized stock return volatility. *The Journal of Financial Economics*, 61, 43--76.
- ANDERSEN, T. G., BOLLERSLEV, T., DIEBOLD, F. X. & LABYS, P. 2003. Modeling and Forecasting Realized Volatility. *Econometrica*, 71, 579-625.
- BAI, J. & NG, S. 2002. Determining the number of factors in approximate factor models. *Econometrica*, 70, 191--221.
- BAILLIE, R. T., BOLLERSLEV, T. & MIKKELSEN, H. O. 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74, 3--30.
- BARIGOZZI, M. & HALLIN, M. 2015. Generalised Dynamic Factor Models and Volatilities: Estimation and Forecasting. *ECARES Working papers*.
- BARIGOZZI, M. & HALLIN, M. 2016. Generalized dynamic factor models and volatilities: recovering the market volatility shocks. *The Econometrics Journal*, 19, C33-C60.
- BOIVIN, J. & NG, S. 2005. Understanding and Comparing Factor-Based Forecasts. *International Journal of Central Banking*.
- BRINER, B. G. & CONNOR, G. 2008. How much structure is best? A comparison of market model, factor model and unstructured equity covariance matrices. *Journal of Risk*, 10, 3--30.
- CHAMBERLAIN, G. & ROTHCHILD, M. 1983. Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets. *Econometrica*, 51, 1281-1304.
- CHAN, L., KARCESKI, J. & LAKONISHOK, J. 1999. On portfolio optimization: Forecasting covariances and choosing the risk model. *The Review of Financial Studies*, 12, 937--974.
- CORSI, F. 2009. A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*, 7, 174-196.
- DAVIDSON, J. 2004. Moment and memory properties of linear conditional heteroscedasticity models, and a new model. *Journal of Business & Economic Statistics*, 22, 16--29.
- DING, Z., GRANGER, C. W. J. & ENGLE, R. F. 1993. A long memory property of stock market return and a new model. *Journal of Empirical Finance*, 1, 83--106.
- ENGLE, R. F. & LEE, G. 1999. A Permanent and Transitory Component Model of Stock Return Volatility. In: ENGLE, R. & WHITE, H. (eds.) *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive WJ Granger*. Oxford University Press.
- FLEMING, J., KIRBY, C. & OSTDIEK, B. 2001. The economic value of volatility timing. *The Journal of Finance*, 56, 329--352.
- FLEMING, J., KIRBY, C. & OSTDIEK, B. 2003. The economic value of volatility timing using "realized" volatility. *Journal of Financial Economics*, 67, 473--509.
- FORNI, M., HALLIN, M., LIPPI, M. & REICHLIN, L. 2000. The generalized dynamic-factor model: Identification and estimation. *Review of Economics and statistics*, 82, 540-554.
- FORNI, M., HALLIN, M., LIPPI, M. & REICHLIN, L. 2005. The Generalized Dynamic Factor Model: One-Sided Estimation and Forecasting. *Journal of the American Statistical Association*, 100, 830-840.
- GEWEKE, J. & PORTER-HUDAK, S. 1983. The estimation and application of long memory time series models. *Journal of Time Series Analysis*, 4, 221--238.
- HALLIN, M. & LIPPI, M. 2013. Factor models in high-dimensional time series—A time-domain approach. *Stochastic Processes and their Applications*, 123, 2678-2695.
- HALLIN, M. & LISKA, R. 2007. Determining the number of factors in the general dynamic factor model. *Journal of the American Statistical Association*, 102, 603--617.

- HAN, Y. 2006. Asset allocation with a high dimensional latent factor stochastic volatility model. *The Review of Financial Studies*, 19, 237--271.
- HARRIS, R. D. F. & NGUYEN, A. 2013. Long Memory Conditional Volatility and Dynamic Asset Allocation. *International Journal of Forecasting*, 29, 258--273.
- JPMORGAN 1994. RiskMetrics™, Technical Document. JP Morgan.
- LUCIANI, M. 2014. Forecasting with approximate dynamic factor models: The role of non-pervasive shocks. *International Journal of Forecasting*, 30, 20-29.
- LUCIANI, M. & VEREDAS, D. 2015. Estimating and Forecasting Large Panels of Volatilities with Approximate Dynamic Factor Models. *Journal of Forecasting*, 34, 163-176.
- MODIGLIANI, F. & MODIGLIANI, L. 1997. Risk-adjusted performance. *The Journal of Portfolio Management*, 23, 45--54.
- ONATSKI, A. 2010. Determining the Number of Factors from Empirical Distribution of Eigenvalues. *Review of Economics and Statistics*, 92, 1004-1016.
- PONG, S., SHACKLETON, M. B. & TAYLOR, S. J. 2008. Distinguishing short and long memory volatility specifications. *The Econometrics Journal*, 11, 617-637.
- STOCK, J. H. & WATSON, M. W. 2002. Forecasting using principal components from a large number of predictors. *Journal of the American statistical association*, 97, 1167-1179.
- TAYLOR, S. J. 1986. *Modelling financial time series*, Chichester, UK: John Wiley and Sons.
- VAN DER WEIDE, R. 2002. GO-GARCH: a multivariate generalized orthogonal GARCH model. *Journal of Applied Econometrics*, 17, 549-564.
- VILASUSO, J. 2002. Forecasting exchange rate volatility. *Economics Letters*, 76, 59-64.
- ZUMBACH, G. 2006. The Riskmetrics 2006 Methodology. *Technical Report*, RiskMetrics Group.
- ZUMBACH, G. 2011. The statistical properties of the innovations in multivariate ARCH processes in high dimensions. *Quantitative Finance*, 1-16.

Table 1. Summary Statistics for the International Stock and Bond Returns

The table reports descriptive statistics for the weekly returns on 34 international stock and bond indices. The sample period is from 01 January 1988 to 31 December 2013. Means and standard deviations are annualised. The normality test is based on the Jarque-Bera statistic. The table also reports the fractional difference operator, d , estimated using the FIGARCH and Geweke-Porter-Hudak (GPH) tests. The GPH estimators are applied to both squared (Sq.) and absolute (Abs.) returns.

<i>Return series</i>	<i>Mean (%)</i>	<i>SD. (%)</i>	<i>Skew.</i>	<i>Kurt.</i>	<i>Min (%)</i>	<i>Max (%)</i>	<i>Norm.</i>	$\hat{d}_{FIGARCH}$	<i>GPH</i>	
									<i>Sq.</i>	<i>Abs.</i>
<i>Panel A. International stocks</i>										
Australia	7.17	22.09	-1.57	18.12	-34.86	14.52	13466	0.36	0.25	0.48
Austria	5.41	26.63	-1.41	15.64	-38.22	20.94	9475	0.49	0.35	0.54
Belgium	6.55	21.21	-1.10	11.14	-26.88	12.53	4017	0.42	0.32	0.47
Canada	7.15	20.48	-1.06	12.80	-25.92	17.61	5679	0.52	0.28	0.47
Denmark	10.79	21.18	-1.23	11.99	-26.39	13.66	4905	0.40	0.25	0.31
France	7.36	22.03	-0.87	9.47	-27.16	13.76	2534	0.34	0.28	0.42
Germany	7.37	23.85	-0.79	8.18	-26.11	15.00	1657	0.39	0.23	0.37
Hongkong	9.03	24.37	-0.61	6.72	-21.08	13.85	867	0.60	0.36	0.49
Ireland	4.60	25.69	-1.62	17.79	-39.31	16.18	12946	0.47	0.41	0.58
Italy	2.13	25.88	-0.64	7.72	-26.71	19.04	1352	0.38	0.37	0.38
Japan	0.09	21.73	0.02	4.77	-16.02	11.75	177	0.23	0.24	0.31
Mexico	17.67	32.49	-0.34	7.83	-30.20	23.23	1343	0.52	0.14	0.32
Netherland	6.66	21.54	-1.26	14.37	-31.48	14.85	7659	0.41	0.39	0.67
New Zealand	1.68	21.43	-0.65	7.36	-23.06	12.07	1172	0.34	0.22	0.15
Norway	8.74	26.79	-0.84	9.63	-28.54	19.82	2642	0.50	0.30	0.40
Singapore	7.05	25.17	-0.68	13.48	-33.13	23.02	6300	0.72	0.46	0.59
Spain	5.85	24.27	-0.78	8.28	-26.22	13.76	1711	0.47	0.28	0.37
Sweden	10.50	27.12	-0.59	7.52	-25.12	19.05	1233	0.50	0.49	0.54
Switzerland	9.07	19.27	-0.73	10.40	-24.01	13.96	3214	0.19	0.14	0.16
UK	5.24	19.11	-1.02	15.15	-27.73	16.30	8571	0.34	0.23	0.37
US	8.22	16.66	-0.76	9.82	-20.19	11.45	2754	0.43	0.37	0.45
<i>Panel B. International bonds</i>										
Austria	1.38	10.60	-0.06	3.61	-5.85	5.72	22	0.31	0.13	0.31
Belgium	1.40	10.96	-0.06	3.88	-6.64	7.03	44	0.34	0.16	0.40
Canada	2.48	8.55	-0.51	6.39	-8.38	5.34	709	0.35	0.27	0.40
Denmark	1.89	10.73	-0.02	3.78	-5.82	5.67	35	0.41	0.24	0.27
France	2.04	10.51	-0.03	3.40	-4.88	5.79	9	0.35	0.27	0.35
Germany	1.09	10.46	-0.01	3.32	-4.52	5.77	6	0.32	0.28	0.36
Ireland	3.33	12.43	0.12	8.04	-7.93	14.11	1440	0.32	0.32	0.35
Japan	1.07	11.81	0.82	8.03	-6.05	14.30	1579	0.27	0.31	0.44
Netherland	1.01	10.52	-0.03	3.31	-4.82	5.45	6	0.32	0.21	0.35
Sweden	0.98	11.99	-0.19	3.78	-7.85	5.93	43	0.49	0.39	0.35
Switzerland	1.62	12.05	-0.12	5.19	-	6.89	274	0.19	0.02	0.24
UK	0.78	10.23	-0.25	4.95	-7.12	6.48	230	0.20	0.45	0.34
US	1.62	4.34	-0.24	3.91	-2.61	2.06	60	0.33	0.21	0.30

Table 2. The Largest Eigenvalues

The table reports the percentage of variance explained by the i^{th} eigenvalue of the variance-covariance and the spectral density matrix.

<i>No.</i>	<i>Variance covariance matrix</i>			<i>Spectral density matrix</i>		
	<i>Eigen-values</i>	<i>Variance explained</i>	<i>Cumulative variance explained</i>	<i>Eigen-values</i>	<i>Variance explained</i>	<i>Cumulative variance explained</i>
1	15.225	0.448	0.448	15.828	0.466	0.466
2	6.649	0.196	0.643	6.679	0.196	0.662
3	1.380	0.041	0.684	1.797	0.053	0.715
4	1.251	0.037	0.721	1.404	0.041	0.756
5	0.930	0.027	0.748	1.132	0.033	0.789
6	0.739	0.022	0.770	0.931	0.027	0.817
7	0.695	0.020	0.790	0.797	0.023	0.840
8	0.610	0.018	0.808	0.685	0.020	0.860

Table 3. Portfolio Performance of the International Stock and Bond Dataset

The table compares the out-of-sample performance of the optimal portfolios constructed using different covariance matrix estimators. Panel A and Panel B report the results for weekly and monthly rebalanced portfolios, respectively. 1/N is the equally weighted portfolio. The static portfolio is constructed using the constant mean and covariance matrix estimated over the initialisation period. LMOF k refers to the Long Memory Orthogonal Factor volatility model with k factors. For each dynamic volatility timing strategy, the table reports the annualised average return (μ), annualised volatility (σ), Sharp ratio (SR), annualised abnormal return (M2) to the static portfolio, annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and breakeven transaction cost τ_γ (in basis points at the rebalancing frequency). Breakeven transaction costs are only estimated when the performance fees are positive.

	μ (%)	σ (%)	SR	M2 (%)	Δ_1	Δ_5	τ_1	τ_5
Panel A. Weekly rebalancing								
1/N	4.480	13.293	0.036					
Static	4.913	3.958	0.231					
<i>Volatility timing strategies</i>								
EWMA	17.325	16.500	0.808	2.282	1113	593	5	3
GARCH-DCC	10.180	20.824	0.297	0.261	317	-532	0	–
LM-EWMA	14.524	14.683	0.717	1.923	861	456	5	3
CGARCH-DCC	8.073	15.798	0.258	0.107	199	-275	1	–
LMOF2	7.428	4.195	0.817	2.320	250	247	26	25
LMOF4	8.296	4.542	0.946	2.830	336	326	28	27
LMOF5	8.048	4.932	0.821	2.335	309	292	23	22
Panel B. Monthly rebalancing								
1/N	4.409	14.012	0.029					
Static	4.892	3.854	0.231					
<i>Volatility timing strategies</i>								
EWMA	16.261	14.068	0.872	2.467	1045	660	9	6
GARCH-DCC	7.130	13.098	0.239	0.029	145	-184	1	–
LM-EWMA	12.308	11.705	0.710	1.844	680	425	8	5
CGARCH-DCC	6.917	12.486	0.234	0.009	132	-164	1	–
LMOF2	7.452	4.352	0.793	2.165	254	245	49	48
LMOF4	8.311	4.565	0.944	2.748	339	327	55	53
LMOF5	8.118	4.805	0.857	2.412	319	301	47	45

Table 4. Average Portfolio Performance of the International Stock and Bond Dataset with Bootstrap Experiments

The table compares the average out-of-sample performance of the optimal international stock and bond portfolios across a wide range of expected returns. A bootstrap procedure is applied to control for estimation error in expected returns. We generate an artificial sample of 4,000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panel A and Panel B report the results for weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap unconditional means and covariance matrices. For each dynamic strategy, the table reports the annualised average returns (μ), annualised average volatility (σ), average Sharp ratio (SR), p -value (proportion) that the dynamic strategy outperforms the static alternative in terms of SR, abnormal return to the static portfolio (M2), average annualised performance fees (in basis points) Δ_γ that an investor with the constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points at the rebalancing frequency).

	μ (%)	σ (%)	SR	p -value	M2 (%)	Δ_I	Δ_5	τ_I	τ_5
Panel A. Weekly rebalancing									
Static	5.275	3.004	0.426						
<i>Volatility timing strategies</i>									
EWMA	13.255	16.998	0.546	0.605	0.365	646	30	3	0
GARCH-DCC	8.872	21.187	0.251	0.252	-0.520	115	-883	0	–
LM-EWMA	11.790	14.089	0.544	0.621	0.361	548	128	3	1
CGARCH-DCC	7.147	14.499	0.235	0.273	-0.569	77	-370	0	–
LMOF2	6.606	2.334	1.119	1.000	2.091	135	142	15	16
LMOF4	7.233	2.553	1.269	1.000	2.543	197	202	19	19
LMOF5	7.069	2.696	1.141	1.000	2.158	180	184	16	16
Panel B. Monthly rebalancing									
Static	5.240	3.119	0.400						
<i>Volatility timing strategies</i>									
EWMA	13.126	15.370	0.596	0.686	0.621	663	129	6	2
GARCH-DCC	7.428	14.056	0.280	0.422	-0.370	112	-341	1	–
LM-EWMA	10.970	11.715	0.591	0.681	0.604	503	206	6	3
CGARCH-DCC	7.187	11.676	0.315	0.452	-0.257	123	-178	2	–
LMOF2	6.539	2.300	1.108	1.000	2.219	132	142	30	33
LMOF4	7.185	2.455	1.300	1.000	2.821	196	204	38	40
LMOF5	7.082	2.561	1.206	1.000	2.527	186	193	34	35

Table 5. Comparison of the Volatility Timing and Static Strategies Using Different Risk Aversion Coefficients

The table compares the average out-of-sample performance of the static and dynamic strategies using different risk aversion coefficients λ . We use a bootstrap procedure to control for estimation error in expected returns. We generate an artificial sample of 4,000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panel A and Panel B report the results for weekly and monthly rebalanced portfolios, respectively. Static portfolios are constructed using the bootstrap unconditional means and covariance matrices. The table reports the annualised average returns (μ), annualised volatilities (σ), the Sharp ratios (SR), p -values (proportions) that the dynamic strategies outperform the static alternatives in terms of SR, average annualised performance fees (in basis points) Δ_γ that an investor with the constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolios, and breakeven transaction costs τ_γ (in basis points at the rebalancing frequency).

λ	<i>Static</i>			<i>LM-EWMA</i>			<i>LM-EWMA vs. Static</i>					<i>LMOF4</i>			<i>LMOF4 vs. Static</i>				
	μ	σ	SR	μ	σ	SR	p -value	Δ_1	Δ_5	τ_1	τ_5	μ	σ	SR	p -value	Δ_1	Δ_5	τ_1	τ_5
Panel A. Weekly rebalancing																			
1	5.275	3.004	0.426	11.790	14.089	0.544	0.621	548	128	3	1	7.233	2.553	1.269	1.000	197	202	19	19
2	4.632	1.498	0.424	8.007	6.959	0.568	0.660	312	209	4	3	5.605	1.268	1.268	1.000	98	99	19	19
3	4.423	1.003	0.423	6.531	4.704	0.534	0.626	199	152	4	3	5.070	0.849	1.262	1.000	65	65	19	19
4	4.319	0.753	0.425	5.863	3.488	0.536	0.635	148	123	4	3	4.801	0.639	1.258	1.000	48	49	19	19
5	4.253	0.600	0.424	5.593	2.823	0.563	0.662	130	113	4	4	4.639	0.506	1.263	1.000	39	39	19	19
Panel B. Monthly rebalancing																			
1	5.240	3.119	0.400	10.970	11.715	0.591	0.681	503	206	6	3	7.185	2.455	1.300	1.000	196	204	38	40
2	4.625	1.576	0.399	7.298	5.901	0.565	0.654	249	175	6	5	5.587	1.230	1.292	1.000	97	99	38	38
3	4.417	1.043	0.401	6.289	3.939	0.581	0.667	179	146	7	5	5.062	0.821	1.297	1.000	65	66	37	38
4	4.310	0.782	0.399	5.724	3.019	0.567	0.649	137	117	6	6	4.789	0.613	1.289	1.000	48	48	37	38
5	4.249	0.630	0.397	5.419	2.400	0.594	0.681	114	102	7	6	4.635	0.491	1.296	1.000	39	39	37	38

Table 6. Yearly Portfolio Performance of the International Stock and Bond Dataset

The table reports the average yearly performance of the international stock and bond portfolios. A bootstrap procedure is applied to control for estimation error in expected returns. The static portfolios are constructed using the bootstrap expected returns and covariance matrices, while the dynamic factor long memory volatility timing portfolios are constructed based on the bootstrap expected returns and forecasts of the conditional covariance matrix from the LMOF4 model. The table reports the average annualised realised returns (μ), the annualised realised volatilities (σ), the Sharpe ratios (SR), the p -values (proportion) that the dynamic strategies outperform the static strategies in terms of the Sharpe ratio, abnormal return to the static portfolio (M2), average annualised performance fees (in basis points) Δ_γ that an investor with the constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points at the rebalancing frequency).

<i>Year</i>	<i>Static</i>			<i>LMOF4</i>			<i>LMOF4 vs. Static</i>			
	μ (%)	σ (%)	<i>SR</i>	μ (%)	σ (%)	<i>SR</i>	<i>p-value</i>	<i>M2</i> (%)	Δ_I	τ_I
1994	6.41	1.75	1.39	7.72	2.08	1.78	0.80	0.70	131	18
1995	2.11	2.13	-0.88	6.98	2.52	1.19	1.00	4.39	486	68
1996	8.18	1.40	3.00	10.23	2.76	2.26	0.07	-1.01	202	22
1997	9.27	2.55	2.08	14.87	3.02	3.59	0.99	3.87	559	62
1998	3.88	3.63	-0.03	5.45	2.56	0.58	0.80	2.22	161	14
1999	12.07	2.66	3.04	5.92	2.22	0.89	0.00	-5.71	-615	—
2000	5.24	3.07	0.41	6.81	2.35	1.21	0.93	2.48	158	14
2001	4.51	4.08	0.13	7.00	2.30	1.29	0.97	4.72	255	23
2002	-2.91	3.67	-1.88	4.23	2.14	0.09	1.00	7.20	719	62
2003	4.90	4.01	0.23	5.23	2.35	0.52	0.71	1.15	38	3
2004	5.49	2.42	0.64	7.85	2.16	1.81	0.94	2.86	237	17
2005	10.36	2.01	3.16	11.61	2.84	2.67	0.20	-0.96	123	9
2006	7.08	2.06	1.50	8.66	3.03	1.54	0.52	0.10	155	9
2007	4.28	2.85	0.10	7.55	3.50	1.05	0.90	2.67	325	17
2008	-3.65	5.09	-1.48	7.87	3.00	1.31	1.00	14.17	1161	74
2009	9.07	3.53	1.43	6.85	1.76	1.63	0.67	0.73	-217	—
2010	5.45	3.04	0.48	7.28	2.18	1.50	0.87	3.12	185	24
2011	1.27	2.94	-0.93	1.97	2.35	-0.87	0.53	0.13	72	10
2012	5.52	1.79	0.84	6.28	2.20	1.03	0.65	0.33	76	14
2013	7.28	1.99	1.65	4.70	2.50	0.27	0.01	-2.73	-259	—

Table 7. Average Portfolio Performance of the DJIA Dataset with Bootstrap Experiments

The table compares the average out-of-sample performance of the optimal DJIA portfolios across a wide range of expected returns. A bootstrap procedure is applied to control for estimation error in expected returns. We generate an artificial sample of 4,000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panels A, B and C report the results for daily, weekly and monthly rebalanced portfolios, respectively. Static portfolios are constructed using the bootstrap unconditional means and covariance matrices. For each dynamic strategy, the table reports the annualised average returns (μ), annualised average volatility (σ), average Sharp ratio (SR), p -value (proportion) that the dynamic strategy outperforms the static alternative in terms of SR, abnormal return to the static portfolio (M2), average annualised performance fees (in basis points) Δ_γ that an investor with the constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points at the rebalancing frequency).

	μ (%)	σ (%)	SR	P - value	M2 (%)	Δ_1	Δ_5	τ_1	τ_5
Panel A. Daily rebalancing									
Static	4.204	1.223	0.169						
<i>Volatility timing strategies</i>									
EWMA	8.119	15.236	0.269	0.687	0.131	274	-195	1	–
GARCH-DCC	4.873	3.517	0.246	0.701	0.101	61	39	2	2
LM-EWMA	5.512	6.058	0.247	0.665	0.104	113	41	2	1
CGARCH-DCC	4.955	3.999	0.237	0.660	0.087	68	38	2	1
FIGARCH-DCC	5.273	13.820	0.130	0.468	-0.053	-10	-481	–	–
LMOF1	4.905	3.379	0.264	0.759	0.120	65	45	4	3
LMOF2	4.966	3.472	0.273	0.800	0.130	71	50	4	3
Panel B. Weekly rebalancing									
Static	4.200	1.167	0.171						
<i>Volatility timing strategies</i>									
EWMA	5.773	7.073	0.251	0.687	0.100	133	33	5	1
GARCH-DCC	4.743	3.503	0.213	0.604	0.052	49	27	5	3
LM-EWMA	5.404	5.734	0.245	0.660	0.091	104	40	4	2
CGARCH-DCC	4.748	3.979	0.190	0.551	0.021	47	18	5	2
FIGARCH-DCC	5.027	7.736	0.132	0.433	-0.056	52	-71	8	–
LMOF1	4.855	3.488	0.245	0.735	0.088	60	38	5	3
LMOF2	4.936	3.572	0.260	0.770	0.105	68	44	9	6
Panel C. Monthly rebalancing									
Static	4.175	1.080	0.170						
<i>Volatility timing strategies</i>									
EWMA	6.051	7.134	0.285	0.733	0.135	162	57	12	5
GARCH-DCC	4.291	2.923	0.104	0.369	-0.063	8	-8	3	–
LM-EWMA	5.266	4.776	0.263	0.710	0.107	98	52	13	7
CGARCH-DCC	4.420	3.602	0.120	0.396	-0.047	18	-6	6	–
FIGARCH-DCC	5.000	6.015	0.176	0.509	0.014	64	-15	8	–
LMOF1	4.646	3.241	0.203	0.578	0.039	42	23	17	10
LMOF2	4.687	3.311	0.209	0.598	0.046	46	26	18	10

Table 8. Comparison of the LMOF and Other Static Factor Models

The table compares the out-of-sample performance of portfolios constructed from different static factor models. We use a bootstrap procedure to control for estimation error in expected returns. We generate an artificial sample of 4,000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. The Factor k is the traditional factor model with k factors, estimated with the Principle Component Analysis. The SMOF k and LMOF k correspond to the conditional factor short memory EWMA and the conditional factor long memory LMOF models with k factors, respectively. For each dynamic strategy, the table reports the annualised average returns (μ), annualised average volatility (σ), average Sharp ratio (SR), p -value (proportion) that the dynamic strategy outperforms the static alternative in terms of SR, abnormal return to the static portfolio (M2), average annualised performance fees (in basis points) Δ_γ that an investor with the constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points at the rebalancing frequency).

	μ (%)	σ (%)	SR	p -value	M2 (%)	Δ_1	Δ_5	τ_1	τ_5
Panel A. The international stock and bond dataset									
<i>Weekly rebalancing</i>									
Static	5.275	3.004	0.426						
Factor4	6.756	2.528	1.097	1.000	2.025	149	155	69	72
SMOF4	7.275	2.645	1.242	1.000	2.460	201	205	38	39
LMOF4	7.233	2.553	1.269	1.000	2.543	197	202	19	19
<i>Monthly rebalancing</i>									
Static	5.240	3.119	0.400						
Factor4	6.750	2.306	1.199	1.000	2.506	153	163	155	164
SMOF4	7.268	2.569	1.276	1.000	2.745	204	211	67	70
LMOF4	7.185	2.455	1.300	1.000	2.821	196	204	38	40
Panel B. The DJIA dataset									
<i>Daily rebalancing</i>									
Static	4.204	1.223	0.169						
Factor2	4.358	2.454	0.144	0.398	-0.032	13	4	12	3
SMOF2	4.976	3.594	0.267	0.786	0.124	71	48	9	6
LMOF2	4.966	3.472	0.273	0.800	0.130	71	50	4	3
<i>Weekly rebalancing</i>									
Static	4.200	1.167	0.171						
Factor2	4.363	2.440	0.150	0.411	-0.029	14	5	51	17
SMOF2	4.966	3.776	0.254	0.764	0.099	70	44	14	9
LMOF2	4.936	3.572	0.260	0.770	0.105	68	44	9	6
<i>Monthly rebalancing</i>									
Static	4.175	1.080	0.170						
Factor2	4.317	2.220	0.145	0.400	-0.029	12	5	91	35
SMOF2	4.786	3.671	0.216	0.618	0.054	55	29	19	10
LMOF2	4.687	3.311	0.209	0.598	0.046	46	26	18	10

Table 9. Comparison of the LMOF and the DFGARCH models

The table compares the out-of-sample performance of portfolios constructed from different factor models. We use a bootstrap procedure to control for estimation error in expected returns. We generate an artificial sample of 4,000 observations by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. The LMOF k correspond to the conditional factor long memory LMOF models with k factors. The DFGARCH(s,d) is the Dynamic Factor GARCH model of s static factors and d dynamic factors. For each dynamic strategy, the table reports the annualised average returns (μ), annualised average volatility (σ), average Sharp ratio (SR), p -value (proportion) that the dynamic strategy outperforms the static alternative in terms of SR, abnormal return to the static portfolio (M2), average annualised performance fees (in basis points) Δ_γ that an investor with the constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points at the rebalancing frequency).

	μ (%)	σ (%)	SR	p -value	M2 (%)	Δ_1	Δ_5	τ_1	τ_5
Panel A. The international stock and bond dataset									
<i>Weekly rebalancing</i>									
Static	5.275	3.004	0.426						
DFGARCH(4,1)	8.141	8.482	0.490	0.955	0.193	255	126	32	16
DFGARCH(4,4)	7.565	4.570	0.790	0.839	1.097	223	197	2	2
LMOF4	7.233	2.553	1.269	1.000	2.543	197	202	19	19
<i>Monthly rebalancing</i>									
Static	5.240	3.119	0.400						
DFGARCH(4,1)	7.966	9.140	0.436	0.860	0.115	235	81	73	26
DFGARCH(4,4)	7.287	3.719	0.886	0.906	1.522	203	194	7	6
LMOF4	7.185	2.455	1.300	1.000	2.821	196	204	38	40
Panel B. The DJIA dataset									
<i>Daily rebalancing</i>									
Static	4.204	1.223	0.169						
DFGARCH(2,1)	4.491	3.077	0.162	0.480	-0.011	25	9	2	1
DFGARCH(2,2)	4.547	3.093	0.176	0.521	0.006	30	14	2	1
LMOF2	4.966	3.472	0.273	0.800	0.130	71	50	4	3
<i>Weekly rebalancing</i>									
Static	4.200	1.167	0.171						
DFGARCH(2,1)	4.523	3.104	0.172	0.498	-0.004	28	11	5	2
DFGARCH(2,2)	4.593	3.100	0.191	0.537	0.018	35	18	6	3
LMOF2	4.936	3.572	0.260	0.770	0.105	68	44	9	6
<i>Monthly rebalancing</i>									
Static	4.175	1.080	0.170						
DFGARCH(2,1)	4.349	2.789	0.132	0.396	-0.043	14	0	8	0
DFGARCH(2,2)	4.428	2.817	0.154	0.429	-0.018	22	8	12	4
LMOF2	4.687	3.311	0.209	0.598	0.046	46	26	18	10

Figure 1. Sensitivity to Estimation Windows: Sharpe Ratios of the Dynamic Portfolios. The average Sharpe ratios of the optimal portfolios constructed from different volatility models are estimated with different estimation windows. Bootstrapped expected returns are employed to account for estimation error. The estimation windows correspond to 4, 6, 8 and 10 years of weekly data for the international stock and bond dataset and to 2, 4, 6, and 8 years of daily data for the DJIA dataset.

Figure 2. The Sharpe Ratios of the LMOF portfolios with Constant and Time-Varying Number of Factors – the International Stock and Bond Dataset.

The figure plots the realised Sharpe ratios for 1,000 trials of the bootstrap experiment for the two portfolios. Each dot represents a separate trial, plotting the realised Sharpe ratios for both the LMOF portfolios with constant and time varying number of factors.